Data Intensive Super Computing at Los Alamos National Laboratory

Hsing-bung Chen, HPC-5; Gary Grider, HPC-DO; Katherine Nystrom, INST-OFF

We discuss here a new big memory footprint machine, built as the computing facility for the IS&T Data Exploratory. The IS&T Data Exploratory is a unique DISC facility designed to explore the challenges of working with large datasets, and provides an open computing facility where researchers can experiment with data intensive methods and research. This facility will give us an opportunity to solve some large-scale graph problems and meet graph computational challenges.

We have utilized unclassified and decommissioned supercomputers and built a new big memory footprint machine (IBMQPI-GraphMach) as the computing facility for the Information Science and Technology (IS&T) Data Exploratory (Fig. 1). The IS&T Data Exploratory is a unique Data Intensive Super Computing (DISC) facility designed to explore the challenges of working with large datasets. Sensors, internet packet filters, telescopes, and satellites all generate massive amounts of data. To solve problems in areas like energy security, biosecurity, and cosmology we need effective ways to analyze the data from these sources. The Data Exploratory provides an open computing facility where researchers can experiment with data intensive computing methods.

The current DISC machine has 720 nodes, each with 8–16 GB of memory and 256 GB to 4 TB of SATA/SSD (serial advanced technology attachment/solid state device) disk space. The 720 nodes will be divided into four segments, each having different usage plans. The first segment (YDS) has 16 GB of memory and 4 TB of disk space on each node and is connected to the Yellow network. The second segment (TDS) is dedicated to the Turquoise network and has 8 GB of memory and 4 TB of disk space on each node. YDS and TDS machines have about two petabytes (2000 TB) of disk space available for data crunching and analysis. The third segment (SDS), which will be used for non-volatile memory and storage system research, has 120 nodes. Each SDS node has one 128-GB SSD and one 120-GB IDE disk drive. The remaining 120 nodes are used for system testing and experimentation, such as the software development and testing of the LLNL Scalable Checkpoint and Restart (SCR) and the LANL Parallel Log File System (PLFS), and the CEPH (an open source distributed file system) file and BTRFS (B-tree file system) file systems testing.

The IBMQPI-GraphMach machine consists of four IBM 3850X5 servers. Each IBM server has 32 CPU cores and 256 GB of main memory. We are using all four IBM 3850X5 servers and setting up a non-blocking network configuration through external Intel QPI (QuickPath Interconnect) links. With the external QPI interconnect we are able to make four IBM servers work as a single big memory footprint machine that comes with 128 CPU cores and 1024 GB (1 TB) of shared memory, and runs a single Linux OS image. This will give us a new opportunity to solve some large-scale graph problems and meet graph computational challenges.

One of the challenges of large datasets is representing the data in a manageable format that scientists can use and learn from. The IS&T Data Exploratory will include a visualization (Viz) environment and, since these DISC machines are in the Yellow and Turquoise networks, we can use existing Viz facilities, which are available in those environments; in addition, we are working towards a new Viz environment in the Los Alamos Research Park building that could also be used. The goal is to find new ways of using visualization for information applications in order to provide an abstraction of the data, rather than a simulation of it.

Currently we are focusing on, but are not limited to, the following research areas for the use of the DISC machines: (1) DISC systems

For more information, contact Hsing-bung Chen at hbchen@lanl.gov

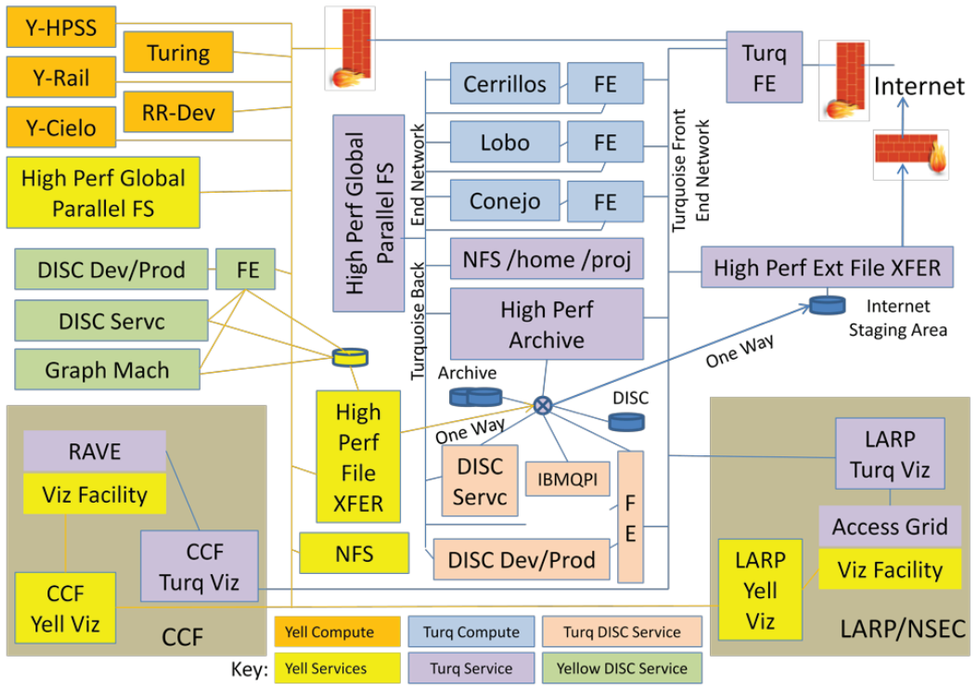

Fig. 1. Open computing environment with Turquoise and Yellow DISC machines.

research, such as bridging the gap of high performance computing and data intensive computing, streaming data, and big data computing, analysis, and processing, (2) graph theory and algorithm research, such as generating triple-store data from the LANL research library collected papers using the Allegrograph semantic database, and studying the federated association of authors and expertise of subject areas, (3) cosmology research, such as conducting data analysis and processing petascale data generated from cosmology simulations, (4) DISC for cyber security, such as security filtering and analysis of network traffic, (5) radioastronomy research, such as processing radio signals for the discovery of radio galaxies, quasars, pulsars, and masers, and (6) optical astronomy research, such as processing optical images for threat reduction and national security. However, we would like to expand into more data-intensive computing areas, such as new algorithms that can scale to search and process massive data sets, distributed file systems, streaming data, a distributed metadata management system to handle complex and heterogeneous data sources, and so forth.

Funding Acknowledgments
Information Science and Technology Institute (ISTI); Information Science and Technology Center (ISTC)